

On the Complexity of Limit Sets of Cellular Automata Associated with Probability Measures

Laurent Boyer^{*1}, Victor Poupet^{**1}, and Guillaume Theyssier^{***2}

¹ LIP (UMR 5668 — CNRS, ENS Lyon, UCB Lyon, INRIA), ENS Lyon, 46 allée d'Italie, 69364 LYON cedex 07 FRANCE

² LAMA (UMR 5127 — CNRS, Université de Savoie), Université de Savoie, Campus Scientifique, 73376 Le Bourget-du-lac cedex FRANCE

Abstract. We study the notion of limit sets of cellular automata associated with probability measures (μ -limit sets). This notion was introduced by P. Kůrka and A. Maass in [1]. It is a refinement of the classical notion of ω -limit sets dealing with the typical long term behavior of cellular automata. It focuses on the words whose probability of appearance does not tend to 0 as time tends to infinity (the persistent words). In this paper, we give a characterization of the persistent language for non sensitive cellular automata associated with Bernoulli measures. We also study the computational complexity of these languages. We show that the persistent language can be non-recursive. But our main result is that the set of quasi-nilpotent cellular automata (those with a single configuration in their μ -limit set) is neither recursively enumerable nor co-recursively enumerable.

1 Introduction

Cellular automata (CA for short) are discrete dynamical systems given by a very simple syntactical definition. They consist of a large collection of identical cells which evolve according to uniform local interactions. Despite the simplicity of the model, they are capable of producing a wide range of different behaviors. One of the main challenges in the field is to give pertinent classifications of these dynamical systems.

There has been a huge amount of attempts in the literature (see [2–4]). Among them, the notion of ω -limit set has received a great interest since the results obtained by K. Čulik *et al.* in [5]. This notion (which comes from classical dynamical systems theory) is an attempt to catch the long term behavior of cellular automata. More precisely, the ω -limit set is the set of configurations that may appear in the evolution after an arbitrarily long time. From a topological point of view, it is also the largest attractor. As shown by J. Kari, ω -limit sets can hold a great complexity since any non-trivial property concerning them is

* laurent.boyer@ens-lyon.fr

** victor.poupet@ens-lyon.fr

*** guillaume.theyssier@univ-savoie.fr

undecidable [6]. Among such properties, the nilpotency is the simplest one: a CA is nilpotent if its ω -limit set is reduced to a single configuration. This property is extremely strong since it implies that all initial configurations lead to the same uniform configuration.

The major drawback of ω -limit sets is that they give the same importance to all configurations. Thus, a negligible set of configurations can influence the ω -limit set of a CA and hide properties of its “typical” behavior.

Recently, P.Kůrka and A. Maass introduced in [1] a notion of limit set associated with a probability measure (μ -limit set). Intuitively, this notion catches the “typical” long term behavior of CA. More precisely, it is defined from the patterns whose probability of appearance doesn’t go to 0 as time goes to infinity. So, as opposed to classical limit sets, it does not deal with what *may* appear in the long term behavior but focuses on what *does* typically appear. This difference makes the μ -limit set more suitable to study some dynamics (see [1]). Moreover, it is a better tool to give theoretical justifications to many phenomena observed experimentally (since experimentations are not exhaustive, they must restrain to “typical” orbits).

In this paper, we mainly study this set from a computational complexity point of view. We first give a new characterization of μ -limit sets associated with Bernoulli measures for any non sensitive CA. Our characterization shows that the μ -limit set does not depend on the measure.

Then we focus on the quasi-nilpotency property: a CA is μ -quasi-nilpotent if its μ -limit set is reduced to a single configuration. One can think that the undecidability behind limit sets disappears as soon as we no longer consider all configuration but only “typical” ones. We show that this is not the case, the Turing degree of the quasi-nilpotency problem is even higher than that of the nilpotency problem: the set of quasi-nilpotent CA is neither recursively enumerable nor co-recursively enumerable. The construction used to obtain this result also allows us to show that some CA have a non recursive μ -limit language.

2 Definitions

Formally, a *one-dimensional CA* \mathcal{A} is a triple $(Q_{\mathcal{A}}, r, \delta_{\mathcal{A}})$, where $Q_{\mathcal{A}}$ is a finite set of states called the *alphabet*, r is the *radius* and $\delta_{\mathcal{A}} : Q_{\mathcal{A}}^{2r+1} \rightarrow Q_{\mathcal{A}}$ is the *local rule*. A *configuration* c describes the state of all cells at a given time: this is a mapping from \mathbb{Z} to $Q_{\mathcal{A}}$. The set of all possible configurations is denoted $Q_{\mathcal{A}}^{\mathbb{Z}}$. For $c \in Q_{\mathcal{A}}^{\mathbb{Z}}$, we will often denote by c_z the value of c at $z \in \mathbb{Z}$.

The local description of the CA induces a global evolution. At every step of the computation, the configuration changes according to the global transition rule $G_{\mathcal{A}} : Q_{\mathcal{A}}^{\mathbb{Z}} \rightarrow Q_{\mathcal{A}}^{\mathbb{Z}}$ induced by the locale rule as follows:

$$G_{\mathcal{A}}(c)_i = \delta_{\mathcal{A}}(c_{i-r} \dots c_i \dots c_{i+r}).$$

In the following, when considering a CA \mathcal{A} , we implicitly refer to the triple $(Q_{\mathcal{A}}, r, \delta_{\mathcal{A}})$, where the same symbol \mathcal{A} denotes both the local and the global mapping.

We denote by $Q_{\mathcal{A}}^* = \bigcup_{n \in \mathbb{N}} Q_{\mathcal{A}}^n$ the set of all finite words over $Q_{\mathcal{A}}$. The *length* of $u = u_1 u_2 \dots u_n$ is $|u| = n$, and, $\forall a \in Q_{\mathcal{A}}$, $|u|_a$ is the number of occurrences of a in u . $\forall 0 < i \leq j \leq |u|$, we also define $u_{[i,j]} = u_i u_{i+1} \dots u_j$ and $c_{[i,j]}$ for $c \in Q_{\mathcal{A}}^{\mathbb{Z}}$ in a similar way. A word u is a *factor* of a word v if there exist i and j such that $u = v_{[i,j]}$.

For every $c \in Q_{\mathcal{A}}^{\mathbb{Z}}$, the *language of c* , denoted by $L(c)$, is defined by

$$L(c) = \{u \in Q_{\mathcal{A}}^* : \exists i \in \mathbb{Z}, u = c_{[i, i+|u|-1]}\}.$$

The language of a subset of $Q_{\mathcal{A}}^{\mathbb{Z}}$ is the union of the languages of its elements.

The *limit set* of a CA \mathcal{A} is given by $\Omega_{\mathcal{A}} = \bigcap_{n \in \mathbb{N}} \mathcal{A}^n(Q_{\mathcal{A}}^{\mathbb{Z}})$. Intuitively, a configuration is in the limit set if and only if it may appear after an arbitrarily long evolution. A CA is said to be *nilpotent* if its limit set is reduced to a single configuration.

For every $u \in Q_{\mathcal{A}}$ and $i \in \mathbb{Z}$ we define the *cylinder* $[u]_i$ as the set of configurations containing the word u in position i :

$$[u]_i = \{c \in Q_{\mathcal{A}}^{\mathbb{Z}} : c_{[i, i+|u|-1]} = u\}.$$

Let \mathcal{A} be any CA and μ any Borel probability measure on $Q_{\mathcal{A}}^{\mathbb{Z}}$ (a measure on the Borel sets, *i.e.* the smallest σ -algebra containing open sets). For any $n \geq 0$, $\mathcal{A}^n \mu$ denotes the probability measure such that for any Borel set $U \subseteq Q_{\mathcal{A}}^{\mathbb{Z}}$ we have $\mathcal{A}^n \mu(U) = \mu(\mathcal{A}^{-n}(U))$. If $Q_{\mathcal{A}} = \{a_1, \dots, a_n\}$ is the working alphabet, a Bernoulli measure μ over $Q_{\mathcal{A}}^{\mathbb{Z}}$ is given by a probability vector (p_1, \dots, p_n) ($0 \leq p_i \leq 1$ and $\sum p_i = 1$) such that, for any word $u \in Q_{\mathcal{A}}^*$ and any $i \in \mathbb{Z}$, $\mu([u]_i) = \prod_{a \in Q_{\mathcal{A}}} p_a^{|u|_a}$. A Bernoulli measure is *complete* (or with full support) if $p_i \neq 0$ for all i .

Definition 1 (Persistent set). *Let \mathcal{A} be any CA and μ be a Bernoulli measure on $Q_{\mathcal{A}}^{\mathbb{Z}}$. A word $u \in Q_{\mathcal{A}}^*$ is a vanishing word for \mathcal{A} and μ if its probability to appear (in a certain position) after n iterations tends to 0 as n grows to infinity. We define the set $L_{\mathcal{R}, \mu}(\mathcal{A})$ of persistent words for \mathcal{A} and μ as the complement of the set of vanishing words for \mathcal{A} and μ : $u \notin L_{\mathcal{R}, \mu}(\mathcal{A}) \iff \lim_{n \rightarrow \infty} \mathcal{A}^n \mu([u]_0) = 0$. Then the μ -persistent set or μ -limit set of \mathcal{A} is the subshift $\Upsilon_{\mu}(\mathcal{A})$ defined by $L_{\mathcal{R}, \mu}(\mathcal{A})$, precisely $\Upsilon_{\mu}(\mathcal{A}) = \{c \in Q_{\mathcal{A}}^{\mathbb{Z}} : L(c) \subseteq L_{\mathcal{R}, \mu}(\mathcal{A})\}$.*

When considering limit sets, the most studied property is certainly the nilpotency. By analogy, we may define the notion of μ -quasi-nilpotency associated with the μ -limit-set.

Definition 2 (Quasi-nilpotency). *Let \mathcal{A} be a CA and μ be any Bernoulli measure over $Q_{\mathcal{A}}^{\mathbb{Z}}$. \mathcal{A} is said to be μ -quasi-nilpotent if $\Upsilon_{\mu}(\mathcal{A})$ is reduced to a single configuration.*

One can verify that a CA \mathcal{A} is μ -quasi-nilpotent if and only if there is some state $q \in Q_{\mathcal{A}}$ such that $L_{\mathcal{R}, \mu}(\mathcal{A}) = q^*$.

Definition 3 (Walls). Let \mathcal{A} be any CA. For any $u \in Q_{\mathcal{A}}^*$, we denote by $[u]_{mid}$ the following set of configurations of $Q_{\mathcal{A}}^{\mathbb{Z}}$:

$$[u]_{mid} = \begin{cases} [u]_{-\frac{|u|}{2}} & \text{if } |u| \text{ is even,} \\ [u]_{-\frac{|u|+1}{2}} & \text{if } |u| \text{ is odd.} \end{cases}$$

A wall for \mathcal{A} is a sequence $\mathcal{W} = (w_n)_{n \geq 0}$ of non empty words of $Q_{\mathcal{A}}^*$ such that:

1. $\forall c \in [w_0]_{mid}, \forall n \geq 1 : \mathcal{A}^n(c) \in [w_n]_{mid}$;
2. the sequence $(|w_n|)_{n \geq 0}$ is non-increasing.

Notice that a wall $\mathcal{W} = (w_n)_{n \geq 0}$ is necessarily ultimately periodic since $[w_0]_{mid}$ contains spatially periodic configurations. The word w_0 is said to be the *foot* of \mathcal{W} . A word is a *foot of wall* for \mathcal{A} if it is the foot of some wall for \mathcal{A} . Any word in the period of the sequence \mathcal{W} will be called a *brick* of \mathcal{W} : formally, w is a brick of \mathcal{W} if there are p, n_0 such that, for all $n \in \mathbb{N}$, $w_{pn+n_0} = w$. A word $w \in Q_{\mathcal{A}}^*$ is a *brick of wall* for \mathcal{A} if it is a brick of some wall of \mathcal{A} .

The following well-known property relates the existence of bricks of wall to the property of sensitivity to initial conditions (see [3] for a proof).

Proposition 1. A CA \mathcal{A} of radius r is sensitive to initial conditions if and only if it has no brick of wall of size r .

The key property behind that proposition is expressed by the following easy-to-prove lemma.

Lemma 1. Let \mathcal{A} be any CA of radius r . If w is the foot of a wall of \mathcal{A} having some brick of size at least r , then for any word $u \in Q_{\mathcal{A}}^*$ there exists a wall of \mathcal{A} whose foot is wuw and which has bricks of size at least $|u|$.

3 Properties of Persistent Sets

It is well-known that the limit set of any CA \mathcal{A} is either reduced to a single configuration or infinite. This fact does not hold with μ -limit sets as shown by the following example. The same example shows a CA whose persistent set does not contain any uniform configuration (the limit set always does).

Example 1. Let \mathcal{A} be the 184 CA in Wolfram's notation. That is, a two states ($Q_{\mathcal{A}} = \{0, 1\}$) one dimensionnal CA of radius 1. Its local rule is given by: $\forall x \in \{0, 1\}, \mathcal{A}(1, 0, x) = \mathcal{A}(x, 1, 1) = 1$ and $\mathcal{A}(0, 0, x) = \mathcal{A}(x, 1, 0) = 0$. It can be seen as a simple model of traffic jam (see [7]).

We first show that for the uniform Bernoulli measure μ_0 , the words 11 and 00 are both vanishing. It can be easily checked that $u = u_0 u_1 \dots u_{2n+1} \in \mathcal{A}^{-n}(00)$ implies that $u_1 u_2 \dots u_{2n+1}$ is a left factor of a well-bracketed string (where 0 "opens" and 1 "closes"). As the proportion of such strings among all words of

length n tends to 0 as n grows to infinity, $\lim_{n \rightarrow \infty} \mathcal{A}^n \mu_0(00) = 0$ and 00 is not persistent. A similar argument shows that 11 is also vanishing.

Because for all n there is at least one word of length n in $L_{\mathcal{Y}, \mu_0}(\mathcal{A})$, and $L_{\mathcal{Y}, \mu_0}(\mathcal{A})$ is stable by factor, and 00 and 11 are not in $L_{\mathcal{Y}, \mu_0}(\mathcal{A})$, we have $L_{\mathcal{Y}, \mu}(184) = (0 + \epsilon)(10)^*(1 + \epsilon)$ and $\Upsilon_{\mu}(184) = \{\omega(01)^\omega, \omega(10)^\omega\}$. \square

We will now give a characterization of the persistent language of non sensitive cellular automata. Before stating the theorem, we need a lemma expressing that for infinitely many steps the preimages of a persistent word must contain any given word at some fixed position.

Lemma 2. *Let \mathcal{A} be any CA of radius r and μ be any complete Bernoulli measure over $Q_{\mathcal{A}}^{\mathbb{Z}}$. Then, for any $w \in Q_{\mathcal{A}}^*$ and $u \in L_{\mathcal{Y}, \mu}(\mathcal{A})$ there are positive integers k_1 and k_2 and a strictly increasing sequence of positive integers $(n_j)_{j \geq 0}$ such that*

$$\forall j \geq 0 : \mathcal{A}^{-n_j}(u) \cap (Q_{\mathcal{A}}^{r n_j - k_1 - |w|} \cdot \{w\} \cdot Q_{\mathcal{A}}^{k_1 + k_2 + |u|} \cdot \{w\} \cdot Q_{\mathcal{A}}^{r n_j - k_2 - |w|}) \neq \emptyset.$$

Proof. Suppose by contradiction that $u \in L_{\mathcal{Y}, \mu}(\mathcal{A})$ does not verify the lemma. Then we have $\forall k \geq 0, \exists n_k \geq 0, \forall n \geq n_k$:

$$\mathcal{A}^{-n}(u) \subseteq Q_{\mathcal{A}}^{n - k|w|} (Q_{\mathcal{A}}^{|w|} \setminus \{w\})^k Q_{\mathcal{A}}^{|u|} (Q_{\mathcal{A}}^{|w|} \setminus \{w\})^k Q_{\mathcal{A}}^{n - k|w|}.$$

Then, for any k and any $n \geq n_k$, we have : $\mathcal{A}^n \mu([u]_0) \leq (1 - \mu([w]_0))^{2k}$. Thus, $\mathcal{A}^n \mu(u) \rightarrow 0$ as $n \rightarrow \infty$ and $u \notin L_{\mathcal{Y}, \mu}(\mathcal{A})$. \square

Theorem 1. *Let \mathcal{A} be a CA which is not sensitive to initial conditions and μ any complete Bernoulli measure. Then $L_{\mathcal{Y}, \mu}(\mathcal{A})$ is exactly the set of bricks of wall for \mathcal{A} .*

Proof. First, consider a brick of wall u for \mathcal{A} . By definition, there exists a word $w \in Q_{\mathcal{A}}^*$ and positive integers n_0 and p such that $\forall c \in [w]_{\text{mid}}$ and $\forall n \geq 0$: $\mathcal{A}^{np+n_0}(c) \in [u]_{\text{mid}}$. Thus $\mathcal{A}^{np+n_0} \mu([u]_0) \geq \mu([w]_0)$ which proves $u \in L_{\mathcal{Y}, \mu}(\mathcal{A})$.

Conversely, let $u \in L_{\mathcal{Y}, \mu}(\mathcal{A})$. By proposition 1, if \mathcal{A} is not sensitive to initial conditions, it has a brick of wall of size at least r (where r is the radius of \mathcal{A}) associated with some wall $\mathcal{W} = (w_n)_{n \geq 0}$. Applying lemma 2 to w_0 , we know there exist positive integers k_1 and k_2 and a strictly increasing sequence of positive integers $(n_j)_{j \geq 0}$ such that

$$\forall j \geq 0 : \mathcal{A}^{-n_j}(u) \cap (Q_{\mathcal{A}}^{r n_j - k_1 - |w_0|} \cdot \{w_0\} \cdot Q_{\mathcal{A}}^{k_1 + k_2 + |u|} \cdot \{w_0\} \cdot Q_{\mathcal{A}}^{r n_j - k_2 - |w_0|}) \neq \emptyset.$$

Since $Q_{\mathcal{A}}^{k_1 + k_2 + |u|}$ is finite, we can extract from $(n_j)_{j \geq 0}$ a sub-sequence $(n_{j_k})_{k \geq 0}$ such that for some $v \in Q_{\mathcal{A}}^{k_1 + k_2 + |u|}$ we have:

$$\forall k \geq 0 : \mathcal{A}^{-n_{j_k}}(u) \cap (Q_{\mathcal{A}}^{r n_{j_k} - k_1 - |w_0|} \cdot \{w_0\} \cdot v \cdot \{w_0\} \cdot Q_{\mathcal{A}}^{r n_{j_k} - k_2 - |w_0|}) \neq \emptyset.$$

By lemma 1, $w_0 v w_0$ is the foot of a wall of \mathcal{A} with a brick of size at least $|v|$. By the above property, we conclude that u is a factor of such a brick of wall. Therefore u is itself a brick of wall of \mathcal{A} . \square

Notice that theorem 1 implies that, for any CA \mathcal{A} which is not sensitive to initial conditions, the set $\mathcal{Y}_\mu(\mathcal{A})$ is the same for any complete Bernoulli measure.

However, there exists some sensitive CA whose μ -persistent set does depend on the Bernoulli measure μ as pointed out by A. Maass and P. Kůrka in [1]: for instance the “just gliders” CA \mathcal{A} is sensitive to initial conditions and such that, for any Bernoulli measure μ , $\mathcal{Y}_\mu(\mathcal{A})$ is reduced to a single configuration if and only if μ gives the same probability to two peculiar letters of $Q_{\mathcal{A}}$.

4 Undecidability Results

This section addresses different decision problems associated with the persistent language of cellular automata. To simplify the statement of the studied problems, we will only consider the uniform measure. Thus, μ will always denote the uniform measure in this section (the working alphabet will be determined by the context). However, all the results extend to complete Bernoulli measures using lemma 2 and theorem 1 from previous section.

Remark 1. In his proof of undecidability of nilpotency [8], J. Kari actually shows that it is undecidable to determine whether a given CA \mathcal{A} with a spreading state (a state s such that $\delta_{\mathcal{A}}(a_1, \dots, a_n) = s$ whenever $s \in \{a_1, \dots, a_n\}$) is nilpotent. Moreover, it follows from theorem 1 that such a CA is μ -quasi-nilpotent for any Bernoulli measure μ (since the only bricks of wall are the words s^n , $n \in \mathbb{N}$). Thus, it is undecidable to determine whether a μ -quasi-nilpotent CA is nilpotent. \square

Theorem 2. *The set of μ -quasi-nilpotent CA is not recursively enumerable.*

Proof. Given a Turing machine M of states Q_M and tape alphabet $\Sigma = \{0, 1, B\}$ working on a semi-infinite tape, we will construct a CA \mathcal{A} of radius 1 that will be quasi-nilpotent if and only if M doesn't halt on the empty input.

The states of \mathcal{A} will be $\{\#\} \cup (S_{\text{simul}} \times S_{\text{signals}})$ where $\#$ is an inalterable state, meaning that if a cell is in this state it will never change to any other state, $S_{\text{simul}} = (Q_M \cup \{-\}) \times \Sigma$ is the set of states needed to simulate the behavior of M (a state $(-, \alpha)$ represents a cell of the tape containing the letter α without the head and a state (q, α) represents that the head is on this cell in state q) and $S_{\text{signals}} = \{-, L, F, R, D\}$ is a set of *signals* whose meaning and behavior will be explained later.

The transition rule of the automaton can be described by the following rules:

- As said earlier, $\#$ states are inalterable. Since the automaton is of radius 1, they act as delimiters or walls, no information can go across them. A finite set of contiguous non- $\#$ cells between two $\#$ states will be referred to as a segment. The length of the segment will be the number of cells between the two $\#$ states.
- At all times, all cells not in the $\#$ state will simulate the behavior of M on their first component. We deal with conflicts (two heads that want to move on a cell for example) in any given way, since we'll see that these have no

impact on what we'll do later (ultimately, we'll only be interested in regular simulations starting on an empty input). If at some point in the computation the head wants to move to a cell in state #, the head is deleted so that the computation cannot end.

- The signal – means that there is in fact no particular signal on the cell.
- If at some point in the computation the final state q_f of M is reached, the cell where this state appears generates a signal F (on its “signal” component).
- The F signal moves towards the left at maximum speed. When it reaches the left border of the segment (#) it turns into an R signal.
- The R signal will move to the right and while doing so it will reset the computation that is held on the first component of the cells it moves through, meaning that it will put the head in its initial state q_0 on the first cell of the segment and put a blank symbol B on every cell of the tape. Since this signal moves at maximum speed, the simulation of M can occur without problems on a clean tape.
- When the R signal meets the right end of the segment it disappears.
- During all this time, the rightmost cell of a segment (any cell that is on the left of a # cell) will generate L signals at every time.
- L signals move to the left at maximum speed. When one of these signals reaches the left border of a segment, it generates a D signal.
- The D signals delete the whole segment by moving to the right while changing all the cells they go through into #. They obviously disappear when they meet a # cell since they can't go any further.

All the signals that we use move at maximum speed (one cell per step) in one of the two available directions. Signals going in opposite directions are not allowed to cross each other, thus, one of the two must disappear. The priority is as follows:

$$L < F < R < D$$

For example, if an R signal is moving to the right (while cleaning the computation) and an L signal is moving to the left, when they meet the R signal keeps moving to the right and the L signal disappears.

Let's assume that M halts in t steps and let's consider the segment of length $2t$ in which there are no signals on any cell, the first cell is in state (q_0, B) and all other cells are in state $(-, B)$. On this segment, the simulation of M starts from a well formed configuration so it will reach the q_f state after t steps and generate an F signal. Meanwhile L signals appear from the right border and move to the left. Because the segment is of length $2t$, the F signal appears on the left of all L signals, so it reaches the origin before all L signals and creates an R signal. This R signal will reset the computation while deleting all L signals. From there a new computation starts that will have enough time to finish again and delete the L signals again. Because the segment is “protected” from any outside interference by the # cells, this cycle will continue forever and no # state will appear on the segment. Because there are only a finite number of possible configurations on the segment the automaton eventually enters a cycle on this non-empty segment.

According to theorem 1 this segment is part of the persistent language so \mathcal{A} is not μ -quasi-nilpotent.

Now we will assume that M doesn't halt and show that any segment of length n disappears after at most $5n$ steps. The proof is based on the observation that we can't delay the apparition of a D signal on the first cell of the segment for more than $4n$ steps.

It's possible that there was already a D signal somewhere on the segment in the initial configuration. In this case, the D signal will cut the segment in two by creating a $\#$ state where it was initially and then delete the right part of the segment. This means that if there is a D signal on a segment in the initial configuration we can focus on a shorter segment on which there is no D initially and let the already present D take care of the rest of the segment.

Therefore we can assume that the segment we are studying doesn't contain any D initially. This means that after at most n steps all original R signals will have disappeared. From there, L signals will start appearing on the right border of the segment and try to proceed to the left (they would arrive at time $2n$). To stop them from reaching the left border and generating a D signal, the only possibility is to generate an R signal on the left border of the segment before the time $2n$. From there, the R signal will reset the configuration of the simulation so that what is computed on the left of this R signal is a normal computation of M on the empty input. Since we have assumed that M doesn't halt, this "well formed" computation will not reach the q_f state. When the R signal reaches the right end of the segment (at time at most $3n$), it disappears and the L signals start moving to the left again. Since the simulation of M doesn't reach the final state no F signal is generated so there's nothing to stop the L signals from reaching the left border, generate a D signal and delete the whole segment. The whole segment is therefore deleted after at most $5n$ steps.

To complete the proof, we need only show that in this case (if M doesn't halt) no other state than $\#$ can appear in a brick of wall. Let's consider a wall $\mathcal{W} = (w_i)_{i \in \mathbb{N}}$. Let's consider the configuration c_{w_0} containing $\#$ states everywhere except on its center where it is the word w_0 . Obviously c_{w_0} is in $[w_0]_{\text{mid}}$ and doesn't contain any segment longer than $|w_0|$ so no segment will survive more than $5|w_0|$ steps, which means that for any $n \geq 5|w_0|$, $\mathcal{A}^n(c_{w_0})$ is the uniform $\#$ configuration, which implies that $w_n \in \#^*$. From theorem 1 we conclude that \mathcal{A} is μ -quasi-nilpotent. \square

Corollary 1. *Given a CA \mathcal{A} and a word w , the property that w is not persistent for \mathcal{A} is not semi-decidable. In other words the set $\{(\mathcal{A}, w) | w \notin L_{\mathcal{T}, \mu}(\mathcal{A})\}$ is not recursively enumerable.*

Proof. We know that a CA is quasi-nilpotent if and only if only one of its states is persistent. If we could semi-decide that a given state is not persistent, then we could use this algorithm on all states in parallel and if the CA is quasi-nilpotent the algorithm would eventually show that all but one states are not persistent, thus showing that the CA is quasi-nilpotent. We would therefore have an algorithm to semi-decide that a CA is quasi-nilpotent, which is in contradiction with theorem 2. \square

Remark 2. The proof above shows that it is also undecidable to determine whether the persistent set is finite or not. Indeed, it is not difficult to check that the persistent set of the constructed CA is either reduced to a single configuration or infinite. \square

Theorem 3. *There exists a CA with a non-recursive persistent language.*

Proof (sketch). It is possible to show this by slightly modifying the CA constructed in the proof of theorem 2. To do so we use another layer in the states so that each regular cell of a segment also has a “memory” containing a tape symbol. The memory of a cell can never be changed (except when the cell becomes $\#$ in which case the memory is lost). Instead of starting from an empty input when the simulation of M is reset by an R signal it’s the memory of each cell that’s written on the tape. This way we can simulate the behavior of M on any input. It is then easy to prove that a segment survives if and only if the memory of its cells corresponds to a word wB^k where $M(w)$ ends using less than $|w| + k$ cells.

If the persistent language of \mathcal{A} is recursive, then the language $\#wB$ such that $M(w)$ halts is also recursive: there is a segment in the persistent language whose memory layer is wB^k , so there is a word of memory $\#wB$ (stability by factor). Therefore if the domain of M is not recursive neither is $L_{\mathcal{Y}, \mu_0}(\mathcal{A})$.

Theorem 4. *The set of μ -quasi-nilpotent CA is not co-recursively enumerable.*

Proof. As with the proof of theorem 2, we will consider a Turing machine M and create a cellular automaton \mathcal{A} of radius 1 that simulates M . \mathcal{A} will be quasi-nilpotent if and only if M halts on the empty input. The idea is that we will again simulate the behavior of M on each segment but now if the simulation doesn’t halt the right $\#$ of the segment will be erased so that the available space for the simulation is increased, and the simulation will start again. If at some point the simulation ends then the segment is erased. This way non-empty segments will remain on the configuration if the machine M doesn’t halt but almost every segment will be deleted if the machine halts.

The construction will be very similar to the one of the proof of Theorem 2. The states of \mathcal{A} are now $\{\#\} \cup (S_{\text{simul}} \times S_{\text{signals}} \times \{0, 1\})$, the new set of signals being $S_{\text{signals}} = \{-, L, R, D, D_L, D_R, C_L, C_R\}$. The added bit doesn’t affect the computation and never changes on a cell. We’ll call it *neutral bit*.

The evolution of the automaton is described as follows:

- The $\#$ state is now “almost” inalterable in the sense that only one particular signal (D) can erase it. We will continue to use the notion of segment (finite set of contiguous cells between two $\#$).
- The simulation of M takes place on each segment as in the previous construction.
- L signals will appear continuously on the right border of a segment and proceed to the left.
- When an L signal meets the $\#$ cell at the left border of the segment it turns into a D signal.

- D signals move to the right. They erase all L signals they meet. If a D signal finds a final state q_f in the simulation of M , it generates two signals D_L and D_R that will erase the segment (turn all cells into $\#$) by propagating to the left and right respectively until they reach the end of the segment. If the D signal doesn't see any q_f state and reaches the right $\#$ of the segment it turns it into a regular cell whose neutral bit is the same as its left neighbor and creates two signals C_L and C_R on the cells next to where the $\#$ cell was.
- The C_L and C_R signals move to the left and to the right respectively. Their function is to clear the segment so that a fresh simulation of M can start back from the beginning. Both signals will erase any signal they come across, C_L having the priority over C_R . When the C_L signal reaches the beginning of the segment it turns into an R signal. When the C_R signal reaches the end of the segment it disappears.
- The R signal moves to the right and resets the simulation as it moves as in the previous proof. It also erases all L signals.

Proving the theorem from this construction will now be similar to the proof of theorem 2. The $\#$ states can only be deleted by a signal that comes from their left so if two segments merge it's because the merging signal came from the leftmost of the two segments while the rightmost one can do nothing to prevent it. We'll say that the the left segment *invades* the right one.

Let's see what happens if M doesn't halt on empty input. In that case a "normal" simulation of M will never reach the q_f state so the D_L and D_R signals should never appear. On any segment where there are initially no signals and no simulation of M going on L signals will appear, reach the left border, and generate an R signal that will start a new correct simulation. This simulation will not end so the segment will eventually invade the one on its right and when doing so C_L and C_R signals will appear to clean the segment and a new correct simulation will again take place on the wider segment, etc. Since no matter how wide the segment is the simulation will never end the segment will never disappear. It is also possible that the segment we have considered is eventually invaded but when the invasion occurs C_L and C_R signals appear that will clean the wider segment and ensure that the new simulation that takes place on this segment is also correct so again there's no risk that the segment disappears.

In other words, if M doesn't halt, any "inactive" segment on the initial configuration will grow and survive forever (the cells that were initially on this segment will never become $\#$). Let s be such an "inactive" segment of length $2k + 1$ including the border $\#$, then for all $n \in \mathbb{N}$ and all $w \in Q_{\mathcal{A}}^n$,

$$wsw \in \bigcup_{q \in Q_{\mathcal{A}} \setminus \{\#\}} A^{-n-k}(q)$$

This means that $\sum_{q \in Q_{\mathcal{A}} \setminus \{\#\}} \mathcal{A}^n \mu(q) \geq \mu([s]_0)$ so at least one of the non- $\#$ states is persistent. By symmetry, changing the neutral bit doesn't change the persistent nature of the state so we have at least two persistent states and \mathcal{A} is not quasi-nilpotent.

Now we have to check that if the machine M halts in t steps then no other state than $\#$ is persistent. Let's consider a segment s of length $l \geq 2t$ on which there is no simulation of M going on and the only signal present is an R signal on the first cell. While this segment is not invaded, it will simply do correct simulations of M , reach the final state in time so that the D signal sees it, the D_L and D_R signals will therefore appear and turn the whole segment into $\#$. This means that such a segment doesn't invade its right neighbor. Moreover, if it happens to be invaded by its left neighbor, the invasion will make C_L and C_R segments appear, which will ensure that on this new segment a new correct simulation starts. The segment will do correct simulations and grow until it's big enough so that a simulation ends. This will happen before all the $\#$ from s 's disappearance have been deleted because there will be enough room to complete a simulation before that so this other segment will also turn to $\#$ before going past the initial boundaries of s . This means that if M terminates there exists a segment such that no matter what happens it will never invade its right neighbor. We'll call such a segment *non-invasive*.

Let's see what happens to a segment such that there is a non-invasive segment on its left at a distance d_l (the distance is taken from the right border of the non-invasive segment to the left border of the considered segment) and a $\#$ on its right at a distance $d_r \geq 2t$ (taken from the right border of the segment).

While the segment is not invaded it will after some time that we can bound easily depending on its length start a correct simulation or be completely deleted (because the L signals cannot be delayed forever). From there it will continuously do simulations and invade its neighbors if the simulations do not halt. Since there is a $\#$ at a distance $d_r \geq 2t$, the simulation will eventually end before this $\#$ symbol is deleted since the segment will be wide enough, so the segment will eventually disappear. The only thing that could delay the disappearance of the segment would be a series of invasion of the segment. However, since there is a non-invasive segment on the left of segment, we know that there is only a limited number of possible invasions so we can bound the time until all possible invasions have occurred. From there, the simulation will start correctly on a segment that will not be invaded and will therefore disappear.

To sum up, we have shown that if M halts, there exists a function $\sigma : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that any segment that has a non-invasive segment on its left at a distance d_l and a $\#$ cells on its right at a distance $d_r \geq 2t$, will disappear after at most $\sigma(d_l, d_r)$ steps. This means that for any $n \geq \sigma(d_l, d_r)$ and any $q \in Q_{\mathcal{A}} \setminus \{\#\}$, any word in $\mathcal{A}^{-n}(q)$ has no non-invasive segment on the cells left of the position $-d_l$ and no two $\#$ symbols on the cells between positions $2t$ and d_r (the first one is the end of the segment, that can possibly be deleted by an already-present D signal). This restriction implies that none of these states is persistent (see lemma 2). \square

Corollary 2. *Given a CA \mathcal{A} and a word w , the property that w is persistent for \mathcal{A} is not semi-decidable. In other words the set $\{(\mathcal{A}, w) | w \in L_{\mathcal{R}, \mu}(\mathcal{A})\}$ is not recursively enumerable.*

5 Conclusion and Perspectives

We proved that the μ -quasi-nilpotency property is neither recursively enumerable nor co-recursively enumerable. In our opinion, such a result has two interesting aspects. First, it deals with a kind of problem rarely considered in the literature: a property of “typical” or random configurations only. We believe that such properties are closer to what experimental observations may capture and therefore that our undecidability results have a stronger meaning to physicists or other scientists concerned with modelling using cellular automata. Second, it gives an example of a “natural” property of cellular automata with a high Turing degree (few examples are known, see [9]).

A natural way to continue the study of the computational complexity of persistent sets would be to try to prove a Rice theorem for μ -limit sets. Any property concerning limit sets is either trivial or undecidable. Is it the same for μ -limit sets?

Another interesting research direction would be to understand better how the probability of appearance of some word can vary with time. More precisely, we left open a very simple question: do we have $L_{\mathcal{R},\mu}(\mathcal{A}) = L_{\mathcal{R},\mu}(\mathcal{A}^t)$ for any CA \mathcal{A} and any t ?

Finally, we can also consider extensions of our work to a broader class of measures or by raising the dimension. In the latter case, the notion of wall does not play the same role (a finite pattern does not cut a bi-dimensional configuration into two disconnected components) and the case of non-sensitive CA is to be reconsidered.

References

1. K urka, P., Maass, A.: Limit Sets of Cellular Automata Associated to Probability Measures. *Journal of Statistical Physics* **100**(5-6) (2000) 1031–1047
2. Wolfram, S.: Universality and complexity in cellular automata. *Physica D* **10** (1984) 1–35
3. K urka, P.: Languages, equicontinuity and attractors in cellular automata. *Ergodic Theory and Dynamical Systems* **17** (1997) 417–433
4. Mazoyer, J., Rapaport, I.: Inducing an Order on Cellular Automata by a Grouping Operation. In: *Symposium on Theoretical Aspects of Computer Science*, Lecture Notes in Computer Science (1998)
5.  ulik, II, K., Pachl, J., Yu, S.: On the limit sets of cellular automata. *SIAM Journal on Computing* **18**(4) (1989) 831–842
6. Kari, J.: Rice’s theorem for the limit sets of cellular automata. *Theoretical Computer Science* **127** (1994) 229–254
7. Nagel, K., Schreckenberg, M.: A cellular automaton model for freeway traffic. *J. Phys.* **2** (1992) 2221–2229
8. Kari, J.: The Nilpotency Problem of One-dimensional Cellular Automata. *SIAM Journal on Computing* **21** (1992) 571–586
9. Sutner, K.: Cellular automata and intermediate degrees. *Theoretical Computer Science* **296** (2003)